


CONNEXIONS



The Interoperability Report

August 1994

Volume 8, No. 8

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Routing Arbiter Architecture..	2
RSVP.....	8
Standards:	
Too much of a good thing?.....	18
Announcements.....	34

ConneXions is published monthly by Interop Company, a division of ZD Expos, 303 Vintage Park Drive, Foster City, California, 94404-1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: connexions@interop.com

Subscription hotline: 1-800-575-5717
or +1-502-493-3217

Copyright © 1994 by Interop Company.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report
and the *ConneXions* logo are registered
trademarks of Interop Company.

ISSN 0894-5926

From the Editor

The Internet continues to grow at a phenomenal rate. Last month Mark Lottor completed another round of the *Internet Domain Survey*. This survey attempts to discover every host on the Internet by doing a complete search of the Domain Name System. The results show more than 3.2 million Internet hosts, a growth of 81% in twelve months. Internet growth is depleting the available address space, as reported in our May 1994 special issue on the next generation of the Internet Protocol (IPng). The Internet Engineering Task Force (IETF) selected the *Simple Internet Protocol Plus* (SIPP) as the basis for IPng at their July 1994 meeting. We will have more details about this decision in an upcoming issue.

In addition to the development and deployment of a new IP, the Internet is undergoing other significant changes, some of which are described in this issue of *ConneXions*. The current NSFNET backbone will eventually be replaced by a system of interconnected network providers. One important component of this new system will be the *Routing Arbiter*. The Routing Arbiter will process topology, connectivity, and routing policy information to create and distribute a stable master routing table. It is described in our first article by Deborah Estrin, Jon Postel and Yakov Rekhter.

In recent years, the Internet has seen the appearance of a number of new and innovative applications. Perhaps the most famous is the *World Wide Web* (WWW). The number of Web servers grew from some 30 in 1993 to over 3,000 in 1994 and new “home pages” are announced every day. While the Web represents a new way to distribute information on the Internet, its use of the underlying infrastructure is not much different from our existing services (e-mail, file transfer and remote login). If, however, the Internet is to be used for real-time interactive audio and video services, new protocols need to be deployed which can “guarantee” certain performance from the network. Our second article, by Bob Braden and Lixia Zhang, describes one such protocol, namely *RSVP: A Resource ReSerVation Protocol*.

This month’s essay by Bob Aiken and John Cavallini discusses the role of standards in the development of network infrastructure. We invite your comments and suggestions for future essays.

Finally another reminder that Customer service and fulfillment for *ConneXions* is now being handled by the Cobb Group in Louisville, Kentucky. You can reach customer service directly at 1-800-575-5717 or 1-502-493-3217. Our editorial address remains the same. We can most easily be reached via e-mail to: connexions@interop.com

Routing Arbiter Architecture

by

Deborah Estrin, University of Southern California and
Information Sciences Institute,

Jon Postel, Information Sciences Institute

and

Yakov Rekhter, T.J. Watson Research Center,
IBM Corporation

Introduction

As the Internet has evolved and grown, a key element has been the provision of stable routing information. An essential element of this has been the *configured networks* information assembled by MERIT and used in the NSFNET Backbone routers. This configured networks information is used to control the dynamic routing information exchange among the NSFNET Backbone and the attached networks.

The Internet is entering a new phase with a new architecture brought about by changes in NSF support. In the new architecture the provision of stable routing information will be much more complex. A new *Routing Arbiter* is to be developed to process the topology, connectivity, and routing policy information to create and distribute a stable master routing table. In addition, the Routing Arbiter will plan for, develop, and deploy new routing services such as multicast, and adaptive alternate path routing.

The new NSF architecture will be comprised of: a *very high speed backbone service* (the vBNS); other *Network Service Providers* (NSPs) which includes both the equivalent of today's regionals, as well as other commercial service providers; *Network Attachment Points* (NAPs) whereby NSPs connect to the vBNS and to one another; stub networks which will connect to their choice of NSPs; and the *Routing Arbiter* (RA) which will manage the routing process for the Internet. In this article we describe the architecture of the Routing Arbiter.

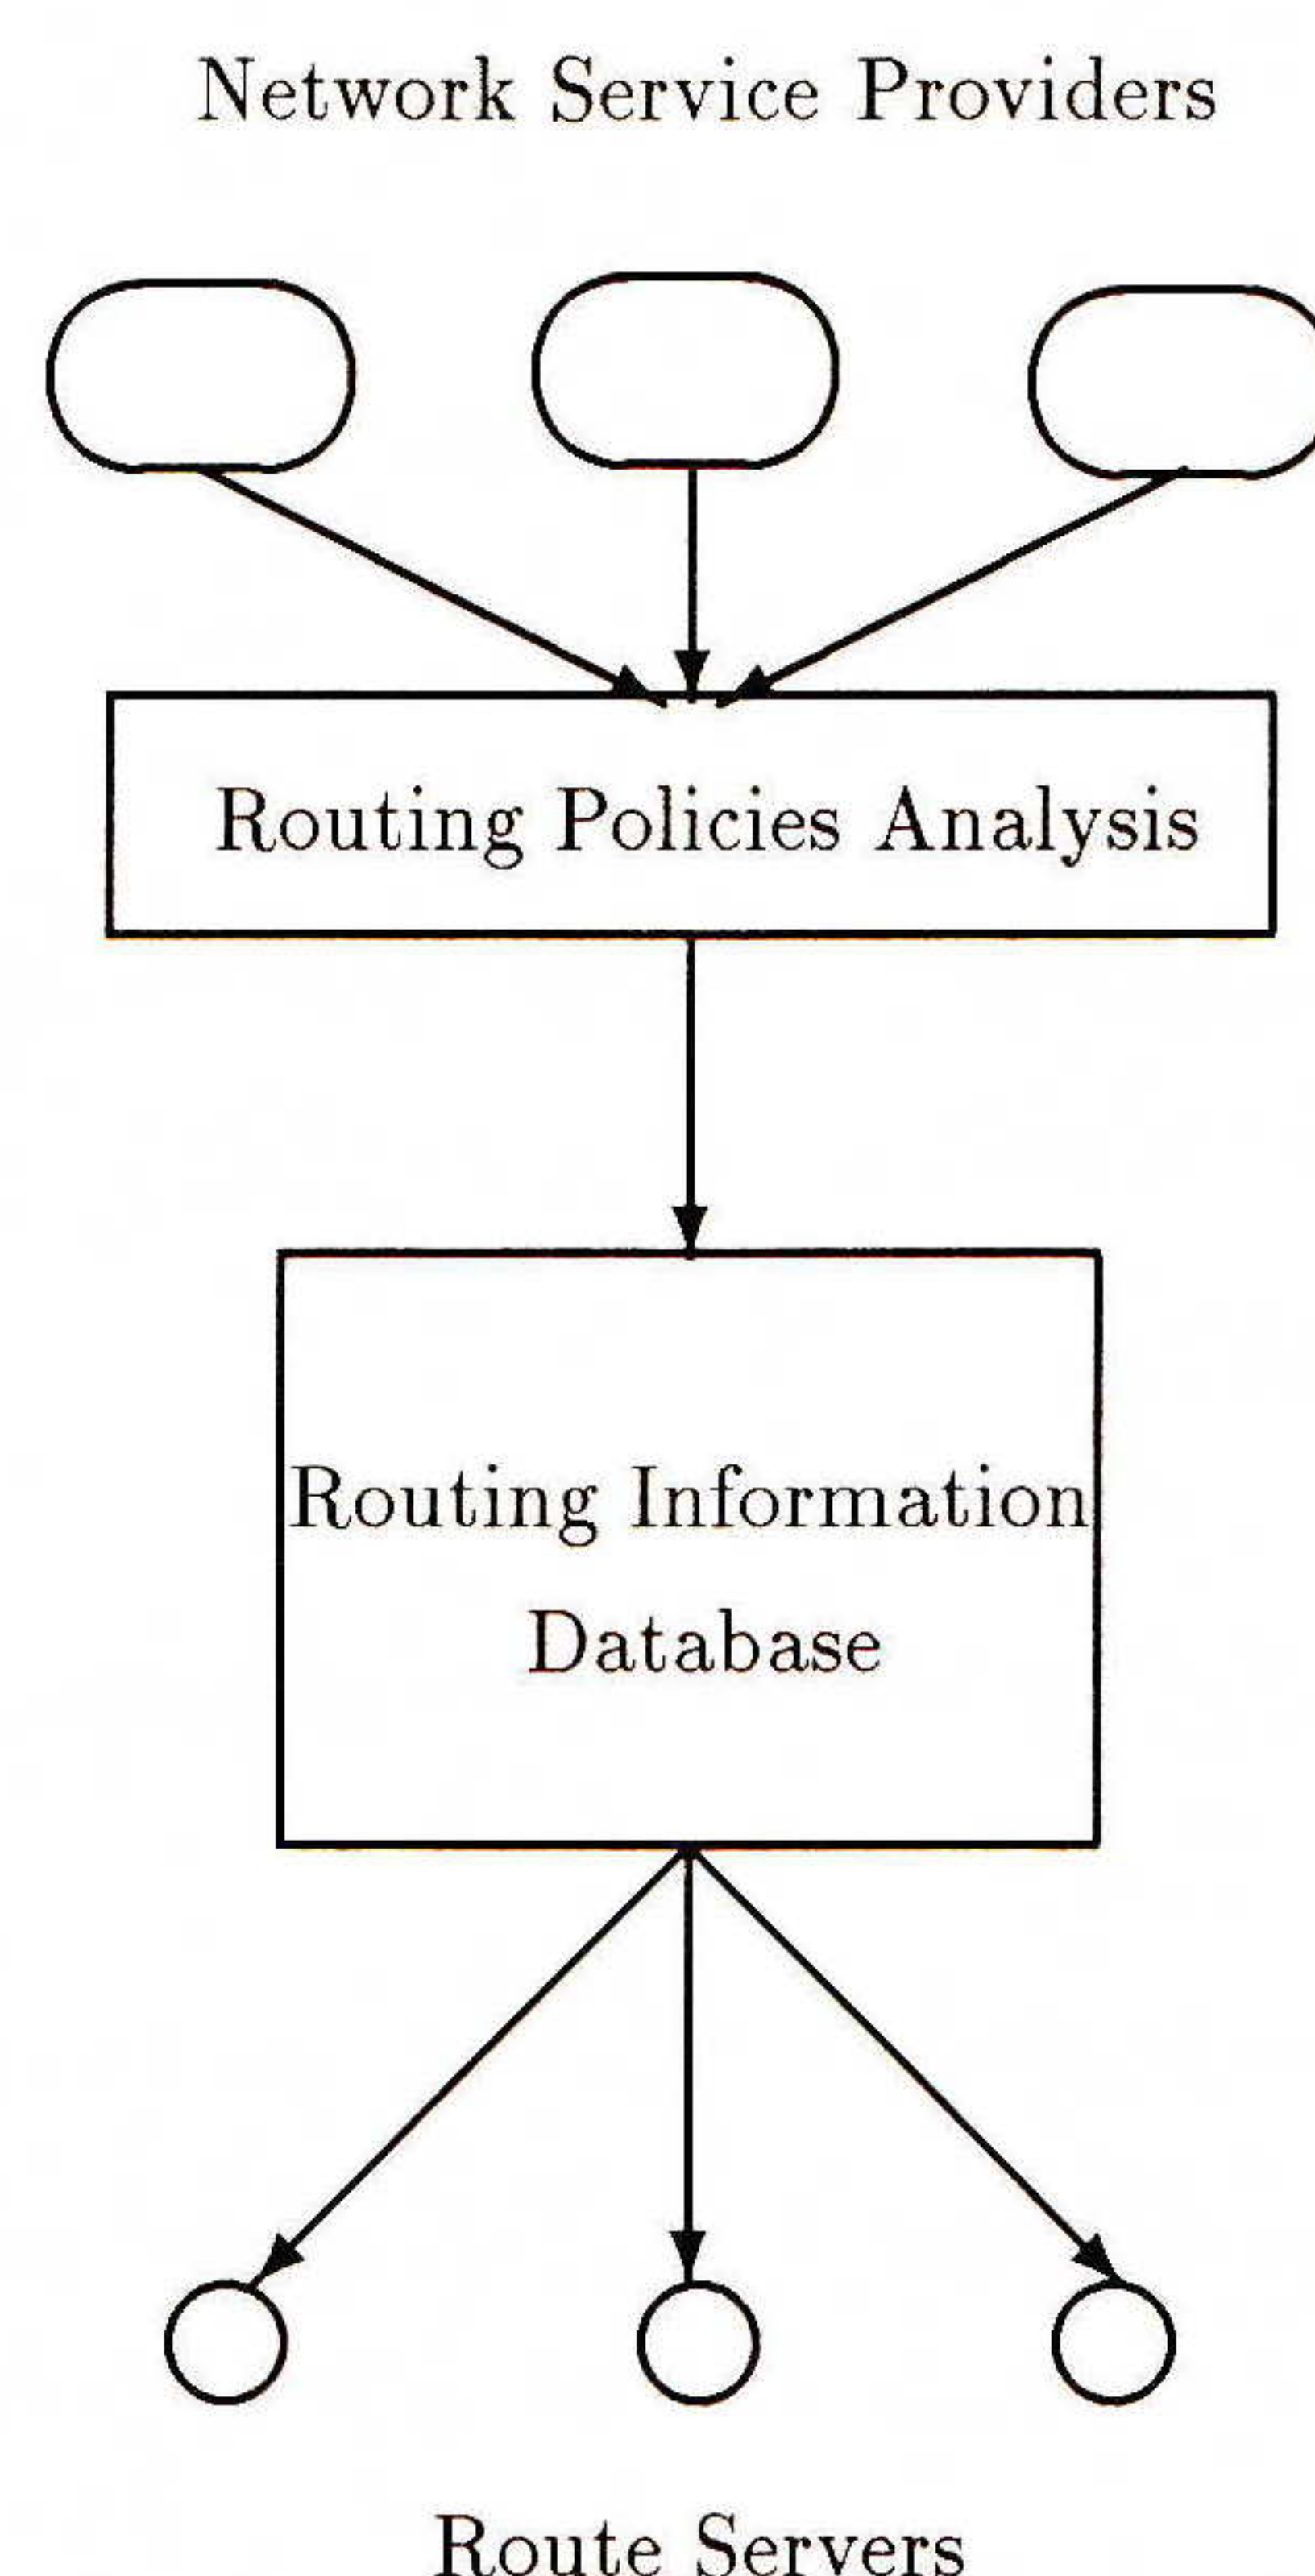


Figure 1: Flow of information between RA components

Proposed Routing Arbiter Architecture

The architecture described in this article is intended to address and solve the problems of the current architecture in the area of routing. Separating the Routing Arbiter into a distinct logical, institutional, as well as operational, component will ensure that: (a) system routing will be analyzable with respect to consistency, connectivity, and configured route characteristics, in the presence of autonomously determined connectivity, configuration, and policy; (b) policies of individual networks will not inadvertently or arbitrarily impact connectivity; and (c) routing will not be used as a means of discriminating or giving preferential treatment to particular NSP(s)—different NSPs serving NSFNET clients (e.g., regional networks) will be given fair and equitable treatment.

The RA architecture consists of the three basic logical components:

- *Routing Registry Data Base (RRDB)*: a logically centralized database of routing information from the Internet's constituent networks, including connectivity, Type of Service, and policy information.
- *Route Servers (RSs)*: one logical server per NAP to control the exchange of routing information among attached networks.
- *Routing Arbiter Network Operations Center (RA-NOC)*: one logically centralized NOC for the Internet.

The flow of information among these components is illustrated in Figure 1. Routing policy and domain-level topology information is submitted by individual NSPs and other clients to the RA. The information is analyzed and subsequently placed in the Routing Registry Data Base. The information stored in the database is used to generate configuration files used by the RSs to control exchange of routing information between attached networks.

Routing Registry Data Base

System-wide connectivity is realized through a composition (synthesis) of the connectivity, TOS, and policy of each participating domain. However, when a domain makes its configuration decisions autonomously, or when such decisions are only coordinated on a bilateral basis between neighboring domains, there is a danger of conflicting or inconsistent routing policies among the domains. Arbitrarily (as opposed to intentionally) incompatible policies can have an unnecessarily adverse impact on connectivity at the global Internet level. The danger is increasing with the proliferation of policy-based routing and the increase in routing policy complexity, which are both consequences of the growth, and consequent heterogeneity, of the Internet.

The function of the RRDB is to maintain the following information for each domain as the input to route construction:

- Domain level connectivity information (i.e., for each domain, what other domains is it connected to directly).
- Aggregation policies (i.e., information about who aggregates routing information for the attached networks).
- Transit Policies (i.e., rules governing use of the network by transit traffic).
- Source and destination policies (i.e., rules governing the use of routes to reach attached sources and destinations).

Routing Arbiter Architecture (*continued*)

This information will be submitted to the RRDB by the Network administrators of vBNS, NAPs, Regionals, and other NSPs, through a combination of manual and automated processes.

In order to implement the RRDB and associated route construction and analysis tools, the Routing Arbiter will define a common *Routing Policy Description Language* (RPDL) that can specify and express routing policy, and topology, information in a globally consistent manner. The RPDL will be used by each domain administration to express the domain's routing policies.

The RRDB will be a vehicle for documenting and sharing routing policy information, analyzing the interaction among the routing policies of different domains and thereby assisting the detection of policy inconsistency and routing coordination in the Internet. The RRDB contents will be used to identify usable routes and diagnose potential problems. The RRDB will be used to generate routing configuration files that will control routing information flow through the RSs attached to each NAP. Users, e.g., network managers, will be able to query the RRDB to conduct their own routing analyses and management functions.

Route Servers

To address some of the architectural problems mentioned earlier the new structure proposed in the NSF solicitation segregates control over the exchange of routing information among the attached networks into a separate entity, the Route Server (RS). RSs address three requirements:

- Promotion of route stability.
- Simplification of routing information exchange between networks attached to a NAP.
- Elimination of bias with respect to different attached NSPs by virtue of being institutionally separate from any NSP.

The policies of attached networks will be checked for inconsistencies and potential conflicts, and analyzed with respect to their impact on the connectivity prior to being used in the routing analysis. Therefore, configuration files constructed from the RRDB, and used by the RS to control routing information flow, are guaranteed to do so without introducing routing inconsistencies and/or instabilities.

To simplify routing information exchange among attached networks, the RS eliminates the need for an attached network to peer directly with all the other networks attached to a NAP—an attached network may peer just with the local RS (i.e., the RS that serves the NAP to which the network is attached). Such simplified peering does not impact the routing information an attached network acquires. Moreover, the RS guarantees that the routing information it delivers to a particular attached network is consistent with routing policies of that network. Consequently, the routing information an RS delivers to different attached networks need not be the same—the RS will “customize” the information according to the routing policies of the individual attached networks.

The elimination of pair-wise peering also simplifies router configuration. Routers can advertise all of the routes obtained from the RS without additional filtering; all of the routing information filtering is taken care of by the RS. In addition, as new routers come online, only the RS needs to be configured with information about this router.

The ability to eliminate pair-wise relationships between the attached networks is not the only simplification an RS can provide. Upon an agreement with the involved attached networks an RS will provide optional aggregation of routing information. The architecture will also provide attached networks with policy-sensitive, and otherwise specialized routes to complement the aggregate routes. This capability will be realized via the the *Source Demand Routing Protocol*, SDRP.

RSs will support exchange of inter-domain routing information via two inter-domain routing protocols, BGP and IDRP. IDRP, will be supported in the “integrated” mode, where it will provide exchange of routing information for both IP and CLNP. BGP will provide exchange of routing information for IP only.

While the protocols supported by RSs are not that different from what is supported by conventional routers, the mode in which these protocols operate in RSs is quite distinct from how these protocols operate in conventional routers. While an RS is intended for use in controlling routing information exchange between the networks attached to a given NAP, the RS is not involved in the actual forwarding of datagrams—routing information exchanged via an RS always includes “third-party” forwarding information. Therefore, RSs allow us to decouple the exchange of routing information from the actual datagram forwarding. Another major difference between an RS and a conventional router is that a conventional router may announce at most a single route to a destination. A RS may announce different routes to the same destination to different attached networks. The only constraint is that to a given attached network an RS may announce at most a single route to a destination. This capability would allow an RS to “customize” the routing information it distributes to individual attached networks to reflect the routing policies of these networks.

While an RS may be used by the attached networks, the proposed architecture does not mandate such use. Individual attached networks may either supplement the information they receive from an RS by direct peering with other attached networks, or may completely bypass services offered by an RS and rely solely on a pairwise peering. However, attached networks that make use of such information must do so at their own discretion; the RRDB and RSs cannot be responsible for routing information obtained “out of band.”

Two Route Server machines will be deployed at each NAP in order to provide a backup capability in case of RS failure. The clients of a particular NAP will connect to the NAP’s RS.

RA-NOC

The *Routing Arbiter Network Operations Center* (RA-NOC) will provide the network operations center functions for the Routing Arbiter. The primary operational concerns are the communication of the current and correct configuration information to all the RSs, and monitoring the correct operation of the RSs. The information gathered by the RA-NOC will be used to alert operators of any unexpected behavior.

The RA-NOC will use SNMP to monitor the RSs. The NOC will also be able to access the RSs via FTP and Telnet. The RSs may also be monitored by the vBNS, NOC, Regional NOCs and other NOCs.

Advanced Routing Services

In addition to managing RSs, the RA-NOC will be responsible for gathering various operational statistics related to the operations of the RA.

Routing Arbiter Architecture (*continued*)

In addition to designing, implementing, and deploying the RA architecture described above, the Routing Arbiter project includes development and introduction of new routing services. Plans are underway to integrate existing and proposed protocols to provide the following services:

- Multicast routing: Multicast routing protocols must run on the network routers themselves and not in the Route Servers per se. However, as part of the RA, we will develop multicast routing capabilities (in collaboration with the existing IDMR IETF working group) and support the planning and coordinated deployment of interoperable, interdomain, multicast technology across the new NSF Internet infrastructure.
- Type of Service routing: TOS routing capabilities will be provided by IDRP and SDRP.
- Indirect provider selection: SDRP will be used to extend the set of dynamically maintained routes available to NAP clients.
- Source Demand Route Construction: SDR routes can be constructed using the RRDB. Users may go directly to the RRDB with route construction requests, or routers may be configured to request particular types of SDR routes from their RSs when needed (i.e., when IDRP computed or existing SDRP routes are not viable).
- Alternate path routing: Reservation traffic and other traffic with special requirements may find that the routes precomputed by IDRP/BGP do not have sufficient resources and will require the capability to request an alternate path. This service will be provided by a combination of RS and router mechanisms.
- IPng: IPng support will be incorporated in a timely fashion.

Participants

The Routing Arbiter will be designed, implemented, and operated through a collaborative effort of Information Sciences Institute (ISI), IBM, and MERIT. ISI and IBM have primary responsibility for routing architecture, software development for RSs and RRDB routing analysis, and introduction of new routing technology (e.g., TOS, multicast). MERIT has primary responsibility for the design and realization of the routing registry data base, the network management architecture, software development for network management and verification tools, and operation of the routing system—including operation of the RRDB and the RA-NOC.

Acknowledgements

We wish to thank Elise Gerich (MERIT) and Kannan Varadhan (ISI) for their extensive comments on this document.

For further reading

- [1] Adams, A., "The MERIT Policy-Routing Configuration System," *ConneXions*, Volume 7, No. 2, February 1993.
- [2] Estrin, D., Rekhter, Y., Hotz, S., "A Unified Approach to Inter-Domain Routing," RFC 1322, May 1992.
- [3] Estrin, D., Rekhter, Y., Hotz, S., "Scalable Inter-Domain Routing Architecture," SIGCOMM 1992, Vol 92, No 4, October 1992.
- [4] Estrin, D., Li, T., Rekhter, Y., "Source Demand Routing Protocol: A Component of The Unified Approach to Inter-Domain Routing," *ConneXions*, Vol. 6, No. 11, November 1992.

- [5] Fuller, V., Li, T., Yu, J., Varadhan, K., "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.
- [6] "Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Network Services Provider for NSFNET and the NRENSM Program," Program Solicitation, National Science Foundation, May 6, 1993.
- [7] "Protocol for Exchange of Inter-domain Routeing Information among Intermediate Systems to Support Forwarding of ISO 8473 PDUs," ISO/IEC/ JTC1/SC6 IS10747.
- [8] Rekhter, Y., Li, T., "An Architecture for IP Address Allocation with CIDR," RFC 1518, September 1993.
- [9] Radia Perlman, *Interconnections: Bridges and Routers*, Addison-Wesley, Reading, Massachusetts, 1992.
- [10] *ConneXions*, Volume 3, No. 8, August 1989, "Special Issue: Internet Routing."
- [11] *ConneXions*, Volume 5, No. 1, January 1991, "Special Issue: Inter-domain Routing."

DEBORAH ESTRIN is currently an Associate Professor of Computer Science at the University of Southern California in Los Angeles where she joined the faculty in 1986. She received her Ph.D. (1985) in Computer Science from the Massachusetts Institute of Technology, her M.S. (1982) from M.I.T. and her B.S. (1980) from U.C. Berkeley. In 1987, Estrin received the National Science Foundation, Presidential Young Investigator Award for her research in network interconnection and security. Over the past ten years Dr. Estrin's research has focused on the design of network and routing protocols for very large, global, networks. Her current research interests are in scalable multicast routing protocols, multicast-oriented reservation setup protocols, inter-domain routing for global internets, adaptive routing to support multimedia applications, and design tools and techniques for developing scalable network protocols. Estrin is a co-PI on the NSF Routing Arbiter project. She co-chairs the Source Demand Routing Working Group of the IETF and is a primary participant in the Inter-Domain Multicast Routing and RSVP working groups. Estrin is a member of the ACM, IEEE, and AAAS. She has served on several panels for the NSF, National Academy of Sciences/CSTB, and Office of Technology Assessment. From 1988 to 1992 Estrin was one of the founding Editors of Wiley's *Journal of Internetworking Research and Experience* and is currently an editor for the ACM/IEEE *Transactions on Networks*. E-mail: estrin@usc.edu

JON POSTEL is the Associate Director for Networking of the High Performance Computing and Communications (HPCC) Division of the Information Sciences Institute of the University of Southern California. Jon has been involved in the development of computer communication protocols and applications from the early days of the ARPANET, and continues to be involved in the development of the Internet. His current interests include multimachine internetwork applications, multimedia conferencing and electronic mail, very large networks, and very high speed communications. Jon is also involved in several Internet infrastructure activities including the Routing Arbiter, the Internet Assigned Numbers Authority, and the RFC Editor. Jon received a BS and MS in Engineering and a PhD in Computer Science from the University of California, Los Angeles. Jon is a member of the ACM and the Internet Society. E-mail: postel@isi.edu

YAKOV REKHTER holds M.S. in Physics from St. Petersburg University, Russia, M.S. in Computer Science from New York University, and Ph.D. in Computer Science from Polytechnic University. He has been with IBM since 1984, and is currently a Research Staff Member and a manager of the Open System Architecture group at T. J. Watson Research Center. Dr. Rekhter was one of the leading architects, as well as the major software developer, of the NSFNET Backbone Phase II. Present activities include work on Classless Inter-Domain Routing (CIDR), work on the Unified Approach to Inter-Domain Routing (under the contract with the National Science Foundation), work on supporting host mobility, and work on supporting IP over Fibre Channel subnetworks. His Internet e-mail address is: yakov@watson.ibm.com

Integrated services

BOB BRADEN holds a BEP from Cornell University and an MS in Physics from Stanford University. During 1962–1986 he held a variety of technical management and research jobs for university computing centers at Carnegie Tech, Stanford, and UCLA and at University College London. He also taught occasional courses in programming, compilers, and operating systems. He entered the network field in 1970, attaching the UCLA IBM 360/91 as a host on the ARPANET. During 1978–1985, he developed a prototype TCP/IP implementation for IBM OS/MVS, before moving to ISI in 1986. He was a charter member of the IAB until this year, and he still chairs the End-to-End Research Group of the Internet Research Task Force. He has been an ACM member since 1958. E-mail: braden@isi.edu

LIXIA ZHANG is a member of research staff at Xerox Palo Alto Research Center. She received her Ph. D. degree in computer science from Massachusetts Institute of Technology. She has been involved in the Internet activities since 1984 and is now a member of Internet Architecture Board. Her research interest includes network architecture and protocols, protocol implementation and performance analysis. She can be reached via e-mail at: lixia@parc.xerox.com

RSVP: A Resource ReSerVation Protocol

by Bob Braden, USC-ISI and Lixia Zhang, Xerox PARC

An increasing fraction of the Internet load is devoted to distributing digital audio and video data streams, in real time. For example, a wide variety of conferences, seminars, and other events are being multicast over the MBONE [5]. These examples represent the leading edge of a large number of future Internet applications; multimedia conferencing, visualization, distributed simulation, remote experimentation (“tele-science”), and virtual reality are all destined to become important in the future.

Broadly speaking, these new applications share common network service requirements. First, IP multicasting [9] is an essential technology for many of them. Without multicasting, IETF multicasts would be hopelessly expensive in bandwidth, for example. There is a great deal of development work in progress on multicast routing, but the importance of multicasting has been established.

The other general requirement of these new applications is for some service “guarantees” from the network. Typically, they need a guarantee of the maximum end-to-end delay or delay “jitter”; we say that such applications need “real-time” service. The present Internet treats all IP packets the same; this common service is known as “best effort” or “as soon as possible” (ASAP). My audio packets and your FTP data packets are put in the same FIFO queue in the routers and thus receive the same ASAP service. However, if your FTP packet arrives late, no great harm is done, while if my audio packet is late, the speech may be garbled. It is therefore somewhat remarkable that audio and video work at all in the current Internet. One reason they do is that Van Jacobson has trained our TCPs to be extremely well-mannered, backing off in the presence of congestion. The other secret is that the applications such as VAT adapt to the variable delays in the Internet. However, there are limitations to the success of these strategies, and in fact real-time applications often do not work very well in the current Internet.

To support this important class of new applications, we must enhance the Internet to provide support for real-time as well as ASAP service, a combination that is known as “integrated service.” This will allow those applications that need it to have a special “quality of service,” while those that don’t can continue to use the ASAP service that is now offered. There has been a flurry of research activities to develop Internet architectural extensions, and in particular a new service model, for integrated service. To provide end-to-end service guarantees for particular packet flows, router and link resources must be reserved for these flows.

An integrated services architecture can be roughly divided into five distinct components [4]:

- (1) A flow specification, or *flowspec*. A flowspec is used to describe the characteristics of the packet stream, or *flow*, sent by the source, as well as the service requirements of the application;
- (2) A *routing protocol* that provides quality unicast and multicast paths;
- (3) A *setup protocol* that allows flows to create and maintain resource reservations;
- (4) An *admission control* algorithm that can maintain the network load at a proper level; and

- (5) A *packet scheduler* in the forwarding path, to assure the committed qualities of service once a flow is accepted by the admission control algorithm.

A more complete description of these components can be found in [4].

RSVP: A setup protocol

RSVP is a setup protocol for resource reservations, designed to provide component (3) of the integrated services architecture [3]. RSVP is an Internet-layer control protocol, analogous to a routing protocol. It does not forward data packets itself; that is still IP's job. Nor does RSVP replace a routing protocol, rather it is designed to coexist with the routing protocols in use today and in the future. RSVP is designed to work with multicast as well as unicast data delivery, and to scale well to large multicast groups. Finally, the setup of resource reservations does *not* force the Internet architecture to adopt end-to-end virtual circuits; rather, RSVP merely introduces "soft state" into the routers, as we see later.

RSVP assumptions

RSVP interfaces to three other components of the integrated services architecture, listed earlier: (1) the flowspec, which is handed to RSVP by an application (or by the host operating system on behalf of an application); (2) the routing protocol(s), which control the forwarding of multicast and unicast data packets, and (3) admission control in each router, which makes an acceptance decision based on the flowspec carried in the reservation messages. RSVP has been designed to make few assumptions about these other components.

For example, RSVP treats a flowspec as a set of uninterpreted bytes of data, passing it to each of the routers along the data delivery path for the particular flow. Each router passes the flowspec to its admission control module, which makes an "accept" or "reject" decision. A resource reservation is established only if all routers along the path accepted the flowspec.

The only assumptions that RSVP makes about the underlying routing protocol(s) are that (1) it provides both unicast and multicast routing, and that (2) a sender to a multicast group is able to reach all group members under normal network conditions (obviously, in the case of a network partition no routing protocol can guarantee reachability). RSVP doesn't assume that a sender to a multicast group is necessarily a member of the group, nor does it assume that the route from a sender to a receiver is the same as the route from the receiver to the sender.




Figure 1: Data distribution to a Multicast Group

The job of RSVP

Although RSVP appears to have a narrow function, its setup job is not trivial. It is designed around multicasting of data, with unicasting as a special case. Figure 1 shows the data flow in a multicast delivery session. Hosts H1 and H2 are sending data to the same multicast group address, while hosts H3, H4, and H5 have joined that group and are receiving the data.

continued on next page

RSVP (*continued*)

The multicast routing protocol in the routers R1, R2, R3, and R4 has set up the multicast delivery trees as shown. The data from H1 and H2 are represented by dark arrows and hashed arrows, respectively.

The job of RSVP is to allow applications to set up resource reservations for these flows. In the process, RSVP should make the most efficient use of network resources, scale well to very large and dynamically-changing multicast groups, and provide flexibility for senders and receivers. To solve this problem, RSVP is being designed with a number of novel features:

- RSVP is receiver-oriented. The receiver of a data flow is responsible for the initiation of the resource reservation for that flow.

This design enables RSVP to accommodate heterogeneous receivers in a multicast group. Specifically, each receiver may reserve a different amount of resources, may receive different data streams sent to the same multicast group, and may “switch channels” from time to time (that is, change which data streams it wishes to receive) without changing its reservation.

- RSVP allows applications to specify how reservations for the same multicast group should be aggregated at the intermediate switches.

This feature results in more efficient utilization of network resources.

- RSVP supports dynamic membership changes and automatically adapts to routing changes by using *soft-state* in the packet switches.

These features enable RSVP to deal gracefully and efficiently with large multicast groups as well as simple point-to-point data flows.

We now discuss these ideas at greater length.

Receiver-initiated reservation

For point-to-point data delivery, one obvious reservation paradigm would have the sender transmit a reservation request towards the receiver, with the switches along the path either admitting or rejecting the flow. This paradigm extends easily to the point-to-multipoint case, by having the sender transmit the reservation request along a multicast routing tree to each of the receivers. This is the technique used by the setup protocol ST-II [8].

However, RSVP’s assumption of multiple heterogeneous receivers and/or multiple senders poses serious challenges that are not addressed by such a source-initiated reservation scheme. These challenges are addressed by RSVP’s *receiver-initiated* design principle. Under this principle, receivers choose the level of service they require, and they are responsible for initiating and keeping a reservation active as long as they want to receive the data.

The receiver knows its own capacity limitations. Furthermore, the receiver is the node that experiences, and thus is directly concerned with, the quality of service experienced by the incoming packets. Additionally, if network charging is deployed in the future, the receiver will probably be the party paying for the requested quality of service. These considerations all lead to the conclusion that receiver initiation of reservation requests is the better choice.

If reservations were initiated by the data source but there were heterogeneous receivers, every receiver would have to send its service request to the source before the source could pass it to the network. The network would then have to figure out, according to the location of individual receivers, how much resource needs to be reserved on each link. For large multicast groups, this approach would scale poorly.

With receiver initiation, each receiver can request resources when it joins the corresponding multicast group. The resource request then travels upstream towards the source(s), but it need travel only as far as the first branch in the multicast distribution tree. Receiver-initiated reservations accommodate heterogeneity among group members while avoiding an implosion of control information at the senders.

Consider, as an extreme example, a cable TV firm that is broadcasting several channels of programs over the Internet. While there are relatively few sources, there may be hundreds or thousands of receivers, each of which can watch only one or a few channels at a time. Using sender-initiated reservations, any individual receiver that wants to switch channels must first send a message to the source. In this case, where there are many receivers and frequent switching between channels, each source would have to accommodate a deluge of change requests. However, this overhead is avoided by letting the receivers ask the network directly for what they need, i.e., using receiver initiated reservations.

Flowspecs and Packet Filters

A resource reservation at a router assigns certain resources (buffers, bandwidth, etc.) to a particular packet flow, as determined by a flowspec. The flowspec does not determine which packets can use the resources, but merely specifies *what amount* of resources are to be reserved. RSVP assumes there is a separate function, called a *packet filter*, which selects those packets that can use the reserved resources. An RSVP reservation request therefore includes both a flowspec and a filter specification ("filter spec"). We assume that filters select packets based upon fields in the headers, such as the IP source and destination addresses as well as port numbers. We expect that production routers will implement packet filters for resource reservation in the same mechanism that is used for route lookups. Thus, packet filters should introduce relatively little extra overhead into the forwarding path.




Figure 2: Filtering by Source

Figures 2 and 3 illustrate RSVP's use of packet filters. In Figure 2, filtering is used to select packets from particular source hosts. Hosts H4 and H5 have included filter specs in their RSVP reservation messages that select data packets sent by H1 only, while host H3 has selected senders H1 and H2. Notice that the effect of the filter specs from H4 and H5 has propagated upstream to router R3, blocking the forwarding of packets from H2 at that point.

RSVP (*continued*)

RSVP provides a choice about how R3 should handle the packets from H2 bound for R4: drop them, or send them without a reservation, i.e., by ASAP service. In Figure 2, they were dropped, avoiding an unneeded load on the R3→R4 link. In either case, a useless resource reservation is avoided on that link.

Filtering may also be more finer-grained, as shown in Figure 3. Here H1 is sending a data stream with two substreams, e.g., two parts of a hierarchically-encoded video signal, represented as two different forms of arrows. Hosts H4 and H5 want both parts, but H3 can only handle the low-resolution signal (dark arrows). Again, RSVP propagates the filter spec upstream to block the unneeded packets, this time on router R1's interface to R2.




Figure 3: Filtering a Substream

Since RSVP distinguishes the filter and the flowspec, it allows the filter to be changed without changing the amount of reserved resources. This enables RSVP to offer several different reservation styles, which we now describe.

Reservation styles

The service requirements of multicast applications dictate how reservation requests from individual receivers should be aggregated inside the network. For example, in an audio conference with several people, there is usually only one person, or at most a few people, talking at any one time, because of the normal dynamics of human conversation. Thus, in many conferencing situations it would be feasible to have all audio sources share the same set of reserved resources on every link. These reservations need only be sufficient for a small number of simultaneous audio streams, no matter how many are in the conference.

Video signals are different: to receive video simultaneously from multiple sources, a separate reservation is required for each of the streams; there are no “silence periods” in video. In Figure 1, for example, R3 must have two separate resource reservations for the distinct video streams going to H3. On the other hand, it needs only a single reservation on the stream to R4, even though that stream is feeding two receivers H4 and H5 (because multicasting sends a single copy of each data packet on the R3→R4 link). Thus, it may be possible to share video reservations among receivers but never among senders, within a particular multicast group.

To support diverse application requirements while making the most efficient use of network resources, RSVP defines different *reservation styles* to indicate how intermediate switches should aggregate reservation requests from receivers in the same multicast group. Currently there are three reservation styles: wildcard-filter, fixed-filter, and dynamic-filter.

We now describe these filter styles; for the sake of brevity we will identify applications only by their multicast address, although in the current Internet context a multicast application may be identified by the IP multicast address plus destination port number.

Wildcard-filter reservations handle the audio case, where a single reservation can be shared by all sources. This is illustrated schematically in Figure 4, which shows a router that might be R4 in Figure 1. The boxes labeled “B” represent the reservations for the two outgoing streams; “B” might be thought of as some particular reserved bandwidth, for example. The shaded lines represent packet flows; each packet arriving on the left is multicast to the two outgoing interfaces on the right. Since the style is wildcard-filter, each interface in Figure 4 has a single reservation that is being shared by packets from both inputs. (Note that some enforcement mechanism is needed to make sure that the aggregate stream does not use more than the reserved amount; this mechanism will be implemented by the traffic control component, not the setup protocol).




Figure 4:
Wildcard filter reservations




Figure 5:
Fixed/Dynamic reservations

Fixed-Filter and Dynamic-Filter style reservations create distinct reservations for different data streams, and are thus appropriate for video signals. As illustrated in Figure 5, each output interface has a separate reservation for each distinct sender. Thus, the Fixed-Filter and Dynamic-Filter styles allow a list of filter specs, each of which can select packets from a particular source; only the packets selected by the filter specs can use the reserved resources. Each filter spec may select particular source hosts as in Figure 2, or substreams as in Figure 3, or both.

A receiver may change its filters at any time to select different senders. However, any change of reservation is subject to admission control in the routers. Suppose that H5 in Figure 2 decided to switch from watching H1 to watching H2. An RSVP filter spec selecting H1 would propagate from H5 towards H1, asking for a reservation. Suppose that there was not sufficient capacity on the R3→R4 link; then the request to “switch channels” would fail. The dynamic-filter style is designed to allow such channel switching without losing, once the initial reservations have been established. The receiver can then “switch channels” by changing the list of sources in the filter at any time during the course of the reservation. Each receiver requests enough bandwidth for the maximum number of incoming streams it can handle at once, and the network reserves enough resources to handle the worst case when all the receivers that requested dynamic filter reservations take input from different sources, even though several receivers may actually tune to the same source(s) from time to time. However, the total amount of dynamic filter reservations made over any link is limited to the amount of bandwidth needed to forward data from all the upstream sources.

continued on next page

RSVP (*continued*)

In summary, having several different reservation styles allows routers to decide how individual reservation requests for the same multicast group can be efficiently merged. So far RSVP has defined three reservation styles; other styles may be identified as new multicast applications, with different needs, are developed. Simulations have shown that RSVP's support for heterogeneous receiver requests and multiple reservation styles can lead to significant reduction in network-wide resource requirements [6]. Preliminary analytical work based on simple network topologies (linear, m-tree, and star) also confirms that, in the asymptotic limit of large multipoint applications, Wildcard-Filter and Dynamic-Filter style reservations can offer substantial savings in resource consumption [7].

"Soft State"

Over the lifetime of a typical real-time application, it is quite possible that new members may join and existing members may leave, and routes may also change due to dynamic status changes at intermediate routers and links. To be able to adjust resource reservations accordingly, transparently to end applications, RSVP keeps *soft-state* in routers, leaving the responsibility for maintaining the reservation to the end hosts. The term "soft-state" was first used by Clark in [1]. In our context, it refers to router state which, when lost, will be automatically reinstated (by RSVP) soon thereafter. Soft-state makes the system self-correcting despite frequent routing changes and occasional service outages.

End hosts must periodically resend the basic RSVP control messages, which carry timeout values that is used by routers to set corresponding timers. These timers are reset whenever a new copy of the same message is received. Whenever a timer expires, the corresponding state is deleted. This timeout-driven deletion prevents resources from being orphaned when a receiver fails to send an explicit tear-down message or the underlying route changes. However, there are also RSVP messages to explicitly delete ("tear down") specific state.

When a route changes, the routing protocol running underneath RSVP will forward future RSVP messages along the new route(s). If the group membership changes, the multicast routing protocol will reflect the change in forwarding the next message. Thus, the state will dynamically adjust to all changes.

Because RSVP control messages are sent periodically, the protocol will tolerate occasional corruption or loss of a few messages. This soft-state approach adds both adaptivity and robustness to RSVP.

RSVP messages

There are two basic RSVP message types, called *Path* and *Resv* (Reservation) messages.

- Each data source sends RSVP Path message downstream to the receivers, following the multicast distribution tree. In Figure 2, the arrows could represent the flow of Path messages as well as the flow of data packets. The Path messages establish or update the "path state" in each router; this is essentially a "trail of breadcrumbs" showing the reverse paths towards the sources. The path state lists a previous hop, an incoming interface, and the IP address of each upstream source.
- Each receiver periodically sends Resv messages upstream towards the senders. In Figure 6, the arrows represent the Resv messages corresponding to Figure 1. The path state is used to route the Resv messages along the reverse data paths.

Each Resv message is processed at each router. The router updates (if necessary) its reservation state. It records:

- (1) The flowspec defining the amount of resources reserved,
- (2) The corresponding filter spec, and
- (3) The reservation style.

After processing, Resv messages are forwarded towards the sources by reversing the paths of Path messages. More specifically, Resv messages of the Wildcard-Filter style are forwarded to all incoming links to the multicast group, and those for the other two styles are forwarded to the previous hops of the sources that are explicitly listed in the filter specs.

Controlling protocol overhead

The advantages of the soft-state approach do not come for free; the periodic refreshing messages add overhead to the protocol operation. To control overhead, RSVP *merges* Path and Resv messages. As a result, in many cases each link will carry no more than one Path message and one Resv message in each direction, for each multicast destination, during each refresh period. For example, Figure 6 shows the Resv messages that will actually flow to make reservations for the data flows of Figure 1. If there were a large number of receivers, the savings on Resv messages on the upstream links would be very large.




Figure 6: Merged Resv messages

At each node, the Resv message that is forwarded upstream carries a flowspec that is the “largest” of the flowspecs that contributed to it. For example, suppose that H4 asked for 64Kbps while H5 asked for 128Kbps in Figure 1; the Resv messages that R4 would send upstream towards H1 and/or H2 would have to ask for 128Kbps. In general, as reservation requests are forwarded along the reverse paths towards the sources, the routers merge the requests for the same multicast group by pruning those that carry a request for reserving a smaller, or equal, amount of resources than some previous request.

The refresh frequencies are controlled by timeout values carried in Path and Resv messages. The larger the timeout value, the less frequently the refresh messages have to be sent. However, there is a tradeoff between the overhead one is willing to tolerate and RSVP’s speed of adapting to dynamic changes. For instance, Resv messages are forwarded according to the path state maintained at intermediate switches, which in turn gets synchronized with the routing protocol state every time a Path message is processed.

RSVP (*continued*)

When a route changes, reservations along the new route (or new route segments) may not be established until a new Path message has been sent (causing the path state to be updated), and a new Resv message has been sent along the new route. However, this pessimistic view assumes that RSVP is strictly separate from routing. If we allow routing to “tap RSVP on the shoulder” when a route changes, then RSVP will be able to adapt locally and immediately to a route change.

Summary

We now summarize the way RSVP operates.

Periodic Path messages are forwarded along the routing trees provided by the routing protocol, leaving behind a “trail of breadcrumbs,” the path state, in the routers.

Receivers send Resv messages upstream along the reverse paths defined by the Path messages, back to the sources. If any router along the way rejects the reservation requested by a Resv message, an RSVP reject message will be sent back to the receiver and the Resv message will be discarded. However, a Resv message propagates only as far as the closest point on the distribution tree where a reservation level greater than or equal to the reservation level being requested has already been made.

When a sender (receiver) wishes to terminate the connection, the sender (receiver) sends out a Path (Resv) teardown message to release the path state or reserved resources. There is no retransmission timer for the teardown message. If the teardown message is lost, the intermediate nodes will simply time out the corresponding state.

Implementation and experimentation

The principal collaborators in RSVP development have been Xerox PARC, MIT, and ISI. These sites have been experimenting with RSVP over DARTnet, the DARPA Research Testbed network. This network is built of T1 circuits linking Sun Sparcstations used as routers. Several prototype implementations have been built, as concepts have evolved. The current version is implemented as a daemon process that runs in user space, in parallel with the unicast and multicast routing daemons.

As we said in the introduction, RSVP is only one component—the setup protocol—within a larger integrated services architecture. The DARTnet experiments use a kernel that incorporates a packet scheduler and an admission control algorithm for the “CSZ” algorithm [2], with IP multicasting. The primary test application is multimedia conferencing. The PARC *nv* packet video program and the ISI *mmcc* conference control program have been modified to invoke RSVP and create resource reservations. Other applications will be converted in the near future.

IETF Working Groups

An RSVP Working Group has been formed in the IETF, to develop a standard RSVP specification, and to coordinate deployment and testing of prototype implementations. This working group, which met for the first time in Seattle in March 1994, is co-chaired by Bob Braden (USC ISI) and Lixia Zhang (Xerox PARC). A parallel working group on Integrated Services will develop the complete service model and architectural framework within which RSVP will be used; it is chaired by Craig Partridge (BBN).

To be added to the mailing lists of these working groups, send a request to:

`rsvp-request@isi.edu` for the RSVP working group
`int-serv-request@isi.edu` for the Integrated Services WG

Conclusions

RSVP is a control protocol to be used by applications to communicate their requirements to the network in a robust and efficient way, independent of specific requirements. RSVP delivers resource reservation requests to the relevant switches, but plays no other role in providing network services. RSVP can communicate the requirements for a wide range of network services, which are provided by other mechanisms. For instance, the flowspec could carry around information about advance reservations (reservations made for a future time) and preemptable reservations (reservations that a receiver is willing to have preempted). RSVP is capable of supporting the delivery of these and other services whenever these network services rely only on state being established at the individual routers along the paths determined by the routing algorithm. Thus, while we have described RSVP as a resource reservation protocol, it can be seen more generally as a router state establishment protocol.

Acknowledgements

Lixia Zhang, Scott Shenker, Deborah Estrin, Dave Clark, Sugih Jamin, Shai Herzog, Steve Berson, Steve Deering, Bob Braden, and Daniel Zappala have all made contributions to the design of RSVP. Jamin, Herzog, and Berson have all worked on experimental and prototype implementations. The original protocol concepts for RSVP arose out of discussions in meetings of the End-to-End Research Group.

RSVP documentation

The latest draft RSVP protocol specification document is available as an Internet Draft. To get a copy of the current draft via e-mail, send a message to `mail-server@nisc.sri.com` and in the body put:

`SEND draft-braden-rsvp-03.ps or .txt`

To get a copy of the latest draft via FTP, connect to `ftp.isi.edu` and do:

```
cd internet-drafts
get draft-braden-rsvp*.ps or .txt
```

References

- [1] D. Clark, "The Design Philosophy of the DARPA Internet Protocols," in the Proceedings of ACM SIGCOMM '88.
- [2] D. Clark, S. Shenker, and L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanisms," Proceedings of SIGCOMM '92.
- [3] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A New Resource ReSerVation Protocol," *IEEE Network*, September 1993.
- [4] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," RFC 1633, June 1994.
- [5] M. Macedonia and D. Brutzman, "MBone Provides Audio and Video Across the Internet," *IEEE Computer*, April 1994.
- [6] D. Mitzel, D. Estrin, S. Shenker, and L. Zhang, "An Architectural Comparison of ST-II and RSVP," in the Proceedings of IEEE INFOCOM '94.
- [7] D. Mitzel and S. Shenker, "Asymptotic Resource Consumption in Multicast Reservation Styles," to appear in the Proceedings of ACM SIGCOMM '94.
- [8] R. G. Herrtwich and L. Delgrossi, "The ST-II Protocol for Multimedia Communication," *ConneXions*, Volume 8, No. 1, January 1994.
- [9] S. Deering, "IP Multicasting," *ConneXions*, Volume 5, No. 3, March 1991.

Essay:
Standards—When is it too much of a good thing?
Will it provide Interoperability for the National Information Infrastructure (NII)?

by Robert J. Aiken and John S. Cavallini,
 U.S. Department of Energy

Introduction

The Administration's goal of a ubiquitous and empowering *National Information Infrastructure* (NII) will require the interconnection and interworking/interoperability of both services and applications. The identification and deployment of appropriate standards on a nationwide basis appears to many to be the obvious solution [5, 17]. In addition to the deployment of an NII, the Administration has strongly committed itself to "Reinventing Government" and has indicated the need for a seamless *Federal Information Infrastructure* (which we will call the *FII* so as not to confuse it with the *Global Information Infrastructure GII*) which will play a pivotal role in achieving the Administration's goal of a reinvented and invigorated government [7, 10]. The Administration, through the *Information Infrastructure Task Force* (IITF) and other forums, will attempt to identify standards, systems, and software products that will be used by the federal government for the purposes of interagency cooperative activities and for interacting with industry and academia [7, 13, 14].

In its drive to create an interworkable NII, the US federal government will identify and specify many standards [17], many formal and also many *de facto*, which the federal agencies will have to implement and deploy; some of these will be successful and therefore aid the development of the NII while others may impede the progress of the NII and subsequently place the US federal government at an international competitive disadvantage. The standards chosen for use by the US Federal Government will be determined and established through a combination of both informal and formal standards processes, mandates, and the sheer purchasing power of the US federal Government through the acquisition of specific technologies or solutions. Various industries, businesses (both national and international), and academia will be directly affected by these choices since they will be required to use them for any business activities they conduct with the US government [11, 12].

Dilemma

A difficult dilemma exists, i.e., should the standards chosen for use by the US government be through a process that attempts to anticipate and precede widespread usage and testing or should the standards and technologies be given time "to be shaken out" by individual agencies, industry, academia and the public relying upon Darwin's theory being applied to technological and standards processes? This essay will discuss these complex issues from the perspective of a user "mission" agency program, whose interactions span the globe and entail both multinational and multiprotocol collaborations with industry, academia, other nations, and its sister agencies. The article will provide some instructional examples of success and failures in adopting and using standards, in addition to drawing some conclusions on the process of choosing and using standards for satisfying agency mission objectives and the broader goals of the Administration.

We intend to show that standards can be either beneficial or a hindrance, depending on how they are developed, adopted and implemented, and that only by using a combination of various standards (be they formal, *de facto* or *de jure*), at different levels of maturity and from different *Standards Development Organizations* (SDOs), can an organization address its requirements in a timely and efficient manner and provide interoperability with other organizations. We also contend that the key to a rich and affordable information and knowledge based economy and society, including the government and the NII, is through the prudent development and use of a diverse and competing set of alternatives and solutions.

The Office of Energy Research's (ER) Office of Scientific Computing (OSC) and its programs, within the Department of Energy (DoE), will provide the backdrop for this treatise. In particular the *Energy Sciences Network* (ESnet), since it is a major component of the National Research and Education Network (NREN) and the NII, has already had to address the conundrum of a multi-standard based environment and its implications with regard to the issues and need for interoperability and its attainment.

Standards: Needs and Perceptions

A discussion of standards, in any computer or information sciences and technology context, is usually accompanied by spirited debate over which protocol or standard is superior. Many currently believe standards to be the only way to achieve, or the panacea for, interoperability which will therefore pave the way to a ubiquitous NII and strong economy [15]. This latter opinion was evidenced at a recent workshop, "Information Infrastructure Forum on Interoperability" (sponsored by The Science, Technology and Public Policy Program and hosted by Annenberg Washington D.C.), where a discussion on interoperability quickly led into a discussion on the adoption of standards as the primary means for achieving interoperability and establishing a healthy information based economy. Most people, in fact, equate the concepts of standards and interoperability. These positions on standards are usually based on the belief, that there always exists or will exist a superior standard that will achieve interoperability on a technical or political (trade and treaty) basis, or some combination of both.

Although standards have been applied to all aspects of the NII, the need for standards is probably evidenced more at the opposite ends of the technology spectrum. At one end of the spectrum, the users desire one user interface and set of tools with which they are familiar, such as a word processor or graphical user interface (GUI), and at the other end of the technology spectrum the service providers desire a small set of physical media interfaces on which they can develop their services and products, such as the Personal Computer Memory Card International Association (PCMCIA) or Narrow band Integrated Services Data Network (N-ISDN).

Issues

Many of the Information Infrastructure Task Force (IITF) Committee on Application and Technology (CAT) working groups, including Health Care and Manufacturing, have issued white papers that identify the need for standards as a prerequisite for making advances on the National Challenges and in other areas. The NII and standards are inextricably intertwined. Some prominent analysts, such as Brian Kahin, see standards not only affecting the success of the NII but also see the NII and its networks as a vehicle "that enables more efficient standards development" [1].

Standards—Too much of a good thing? (*continued*)

The majority of people involved in the development and deployment of technology, in particular those relevant to the NII, would probably agree that some base sets of standards are necessary for interoperability and for the success of the NII. However, when trying to identify which and how many standards are necessary to achieve these goals, many issues arise such as:

- What standards exist and need to be developed and enforced? (e.g., physical layer network standards and application standards such as SONET and Health Level 7, aka HL-7)
- How should standards be chosen? (e.g., what process, what involvement by industry, influence purchasing power of federal government, actions of formal standards bodies, consortia, etc.)
- Who chooses the standards for defense and civilian federal agencies? (e.g., can each agency or organization be allowed to make its own choice or do they have to follow the directives of NIST/OMB/GSA)
- Who are the people actually developing and mandating standards? (e.g., do they have real-life operational experience in the area they are so greatly influencing)
- What are the professional and ethical responsibilities of those persons and organizations who set standards? (e.g., is short term cost benefit or conformity more important than diversity and competition)
- Should the U.S. Government formally require and mandate standards? (e.g., ADA, GOSIP, Clipper Chip, etc.)
- Should multiple standards be allowed to coexist? (e.g., at the network layer are IP and CLNP allowed to coexist?)
- What is the real practical life cycle of a technology and/or standard and how is it phased out or replaced when appropriate?
- How do government purchasing practices (e.g., two FTS2000 network providers for all federal agencies) affect both the standards process and the competitive technology and services based marketplace?
- Should the government use its awesome purchasing power (e.g., FTS2000) to set a procurement induced standard and then seek to create a market for those standards to ensure lower costs to the government and interoperability?

Standards and interoperability

Standards have traditionally been adopted to achieve interoperability and to provide a common base for multiple applications; however, this does not necessarily guarantee interoperability. Some organizations have tried to use X.400 compliant products for their electronic mail only to find that two different vendors compliant X.400 systems do not interoperate since each uses a different portion of the X.400 address for routing purposes. *Narrow band Integrated Services Data Networks* (N-ISDN) is a standard that has been implemented in many regions, yet there has been a lack of inter-vendor ISDN interoperability at a national level. Even different versions of one standard may have interoperability problems such as the loss of functionality between two different versions of the same standard (e.g., X.400 1984 and X.400 1988).

Technology and standards continuum

Many standards have been created based on the wide spread acceptance and use of formerly proprietary standards (e.g., NFS) and/or novel tools or applications, such as *Wide Area Information Services* (WAIS) and the *World Wide Web* (WWW). These latter examples were not specifically designed for the purpose of interoperability and yet they have become *de facto* interoperable standards for their respective applications by satisfying an urgent need. In addition, interoperability can, and has been, achieved by cooperating and consenting organizations based on agreement(s) between the organizations involved to use specific technologies (some are standards based—some are not) for the purpose of interoperating. There are many valid reasons and ways to choose standards, and a variety of forums or methods for developing them. It is clear that the formal standards arena is not, and should not be considered, the only means of achieving interoperability. The realistic approach to attaining interoperability is based on the use of a combination of various formal, *de facto*, *de jure*, consortia, and even proprietary standards, all at different maturity levels in the respective technology and standards life cycles.

When discussing standards and analyzing their possible use, consideration of the technology and standards continuum is important since it will affect the ability of any organization or individual in their attempt to identify, design, adopt, or affect a technology standard, and subsequently any operational services based on that standard. The technology and standards time continuum is composed of many technologies and standards at different maturity and evolution stages. Therefore, at any given discrete point in time, the life cycle of a technology will be intersected and possibly paralleled by the life cycle of a relevant standard associated with that technology, and likely intersect with that technology's and standard's predecessor(s) and successor(s). To further complicate matters, this process is non-linear. The successor to a particular standard or technology may be injected into the continuum from a non-traditional source, and having no known lineage on the continuum it is a newcomer and thereby introduces a bit of chaos into the continuum. In addition, there exists more than one SDO and therefore multiple standards activities parallel, overlap and intersect each other on the technology time continuum. The current standards processes are usually based on the concept of a single formal SDO and a linear technology and standards continuum. Yet reality has proven otherwise, as evidenced by the failure of so many anticipatory standards. The continuum is non-linear and multi-dimensional. The short cycle associated with current computer and telecommunications technologies makes long term forecasting a very challenging endeavor.

As an example, both the OSI and TCP/IP network layer protocols were not originally designed to support the real time resource management necessary for videoconferencing, live simulations, and other demanding real time applications that are becoming popular today [6]. Many sectors, such as the manufacturing sector, the power and utility sector, government Administration, and others, invest a lot of capital in equipment. Their investments have traditionally been made with the expectation that they will usefully serve their requirements for long periods of time and the investments can be gradually depreciated. In addition, many of these large industries operate globally and are critically dependent on computers, telecommunications, and information services to remain competitive.

Standards—Too much of a good thing? (*continued*)

These industries are now facing very challenging issues when deciding which technologies and standards they should employ as part of their strategy since many of their investments are now more tightly coupled to the fast paced evolution cycle of the computer and information areas. The prior reliance by these sectors on the international and slower paced standards is now being augmented and in some cases supplanted by standards from the *de facto* and consortia standards arenas, which at times can seem random and chaotic in contrast to the slower moving pace of the formal standards process. By the time a standard has been identified for use in any arena, its replacement or a competing solution is already in the pipeline and being tested somewhere. The leading edge standard of today is tomorrow's legacy standard and system. The technology standards continuum is comprised of both anticipated and surprise developments. Just as e-mail was the surprise application that arose from the ARPANET testbed there lurks another such application(s) in the near future (e.g., Internet talk radio or packet video?) that may drive us in a direction we had not planned for nor anticipated with respect to standards planning and development. For this reason, exclusively mandated anticipatory standards (e.g., GOSIP/OSI) are not likely to succeed, even when the government tries to create a market for these products through procurement policy. The old standards paradigm, which was based on long drawn out, formal, and exclusive standards processes, no longer seems relevant given today's fast-paced and seemingly chaotic technological advances.

Time intervals

One of the major reasons that standards such as GOSIP, OSI, ADA and others have never achieved their full potential is that standards development takes too long. These standards were overcome by events and overtaken by fast track *de facto* implementations, such as TCP/IP and C/C++. The methods and processes for developing standards, such as OSI and other formal standards, are too long and complex and therefore are no longer a true metronome of the fast paced technology time continuum. It is interesting to note that the growth of the *Internet Engineering Task Force* (IETF), and subsequently its consensus process, has recently strained its capability to resolve major standards issues in a timely fashion (e.g., the problems encountered with IPng—the next generation of IP) and it is now enduring some of the same problems and issues normally associated with the formal standards organizations. Since the formal standards processes are slow it is natural to assume that faster paced responses from other sources (e.g., *de facto*, consortia, etc.) will develop. Given that situation, in addition to the fact that any standard will have versions that may or may not interoperate, coexistence is a reality that must be addressed.

The Federal Government's current attempt to reconcile such coexistence issues (the *Federal Information Requirements Panel*—FIRP), which are a direct result of the natural and non-deterministic technology and standards continuum, through the strategy of adopting new standards (e.g., making TCP/IP part of GOSIP as recommended by the Federal Information Requirements Panel) does not solve the problem since the issue is rooted in the ongoing, competitive, and iterative cycle on the technology and standards continuum, and will therefore be repeated again. Although this attempt to address the coexistence of TCP/IP and OSI was noble, it would have been more effective and efficient for the FIRP to recommend that no one standard or SDO be given policy preference.

Also, the current process of mandating or specifying standards no longer works since there will always be yet another standard or technology in the pipeline that will need to coexist with and eventually replace the existing standards. The standards continuum will always be comprised of the past (e.g., legacy systems, proprietary e-mail, etc.), the present (e.g., TCP/IP and GOSIP, SMTP and X.400, etc.), and the future (e.g., IPng, MIME, etc.), with the time intervals on the continuum remaining short and overlapping one another. A standard whose conceptual design and implementation adheres to the fast-paced technological development and the application-level demand is a *Just In Time Standard* (JITS) whose probability of successful adoption and implementation is very high. Yet, no matter how hard we strive to develop standards, including just in time standards, to allow for evolution and change (e.g., version numbers, etc.) we must realize that any standard will ultimately be overcome, the only question is: how quickly? Given today's fast-paced technological and standards cycles, the times between innovative advances is very short and this means that standards, from multiple SDOs, will need to be developed and replaced to meet these short technology cycles if standards are to be effective or even contribute to solving interoperability issues.

Measured response

The authors believe that some base set of standards is essential to the success of any technological endeavor, especially for the NII, but that these need to be applied in measured doses at the proper time in the technology life cycle continuum. We also believe that the overzealous creation and use of standards, either through formal standards processes or by government purchasing practices, poses the risk of impeding the introduction of necessary new technologies and services for use in the NII and can additionally adversely affect the competitive marketplace, whose healthy existence is essential for the success of the NII and achieving the Administration's goals for national competitiveness. In addition, the "technology and innovation gene pool" of the future is jeopardized when choices are made for short term benefits which can be achieved through other ways, especially when policy mandated anticipatory standards prematurely prune a viable branch from the evolutionary tree of technology. The major advances and health of US technology are primarily due to the rich diversity of competitive techniques, ideas, and solutions that have been fostered in the past. The explosion of the Internet and its information discovery and retrieval tools (e.g., *WAIS*, *Mosaic*, *Gopher*, etc.) is a testimony to this concept. These tools, which all branches of the federal government are using today to reach out to the public, may not have existed had everyone exclusively used the federal government's choice of OSI/GOSIP protocols and products.

The method with which industry and government view and use standards and their processes needs to undergo a paradigm shift. The technological life cycle and the introduction of new ideas is very short (6–18 months) and therefore old standards setting processes and paradigms are no longer valid. Kuhn [8] identified and described the cycle of scientific revolution and its importance in the evolution of science. It is important to recognize that the same metaphor applies to the technological and standards continuum and that we are currently at one of those revolutionary junctions in the continuum with the golden opportunity to affect the necessary paradigm shift with respect to standards and their processes.

Standards—Too much of a good thing? (*continued*)

Further evidence of this juncture point is identified in Clifford Lynch's [9] treatise on libraries and relevant standards, and Brian Kahin's [1] assertion that we are already in the transition to a Knowledge Information Infrastructure that will require a new look at standards and how they affect our ability to function in a knowledge and information based society. Drucker [4] also asserts that we are in the process of transitioning to an information society that will require all of us to dramatically change the way we interact and do business on both a national and international basis.

Much of the current attention to standards and their processes have focused on "fixing" or "re-engineering" the current standards process. The new paradigm should be based on the pragmatic reality that no one standard or SDO will satisfactorily address our requirements in this fast moving technological evolution to the information and knowledge society. Competitive diversity, with a small bit of chaos, is beneficial and encourages alternatives for current and future requirements. Competition among SDOs will ensure that they remain responsive to the needs of their communities and constituents or they will be replaced by more adaptive and responsive SDOs. Also, the concept of one standard satisfying all requirements within a given area is a fallacy. There will always be versions of standards as they evolve in addition to new or alternative standards and technologies being developed to satisfy requirements and needs not addressed by current standards, for whatever reason. Therefore, the new paradigm should also be based on the realization that the support and exploitation of multiple standards and technologies will provide a healthier technological future since any agreed upon standard will have been chosen through the natural selection process from among its competitive peers, and consequently will more likely address the needs of the community affected.

Standards and interoperability in action (examples)

Bearing in mind that the technology and standards cycles utilized by DoE programs have very short life spans that can become eclipsed by newer and usually better technologies, we will attempt to further elucidate the above issues with the following real life examples. The Energy Sciences network (ESnet) is an evolution of the *Magnetic Fusion network* (MFENET), *High Energy Physics network* (HEPnet) and other energy research community based networking activities. HEPnet was predominately based on DECnet and MFENET was based on protocols that were developed in 1974–1975 by the *National Energy Research Supercomputer Center* (NERSC), formerly the *National Magnetic Fusion Energy Computer Center* (NMFECC), to provide remote access to the center's resources. The MFENET protocols were based on a draft version of the DECnet protocols that preceded the first release of a DECnet product. Yet as both DECnet and MFENET evolved, they did so on different paths. This difference in direction, coupled with the need to provide for the network protocols on its supercomputers, left MFENET the task of supporting a non-vendor proprietary protocol stack for a very long time. Clearly, at the time the MFENET was originally designed the center made the correct choice to implement its own protocols since there existed no other viable solution for a satellite based network connecting users and supercomputers at that time. One problem encountered later was that MFENET did not transition to an open standards based network protocol as quickly as was prudent. This was mainly due to an installed user base which was reluctant to change, even to achieve interoperability, and due to the large effort required to ensure a smooth transition from MFENET to a new network architecture.

Therefore, MFENET did not cease to exist until sometime in 1991, thereby requiring staff and effort to support it and to provide for the transition to the current multi-protocol ESnet. The use of pre-standards and/or niche application solutions can and should be used to satisfy time critical requirements, yet they can also create problems if the transition to new standards, whether formal or *de facto*, is slow to follow. However, the coexistence of different protocols in this case satisfied the evolving and varied requirements of the user community and provided the capability for a smooth transition from older solutions.

Evolution

In 1987, the ESnet staff had already decided to implement the TCP/IP protocols as the main basis for their networks, while still pragmatically supporting other protocols where appropriate. In 1989 a peer review of ESnet supported their choice to move to a TCP/IP based protocol suite as a replacement for the MFENET protocols. ESnet chose to acquire commercial routers that would provide the capability for supporting TCP/IP, DECnet IV, OSI/CLNP, X.25 and other protocols. In 1989, a national multiprotocol router based network was a fairly new concept and some were skeptical about ESnet's success. However, ESnet is a customer driven network service and therefore needed to provide both national and international support and access to a variety of sites whose applications required different standards based protocols. Contrary to the popular belief that a single standard is the only means to successfully support a diversified community (e.g., NII), ESnet has demonstrated that the support of multiple standards based services addresses the users needs and provides for the coexistence necessary to allow the natural selection process to occur. This provides the "shake out" time period necessary to fully test and implement alternate solutions in addition to providing the basis for transition from older solutions to new ones. Interoperability is achieved through the support of application based gateways thereby supporting each constituency's current requirements and providing a transition path to future standards and systems. By not exclusively focusing on or choosing a single standard, either the federally mandated GOSIP or the popular *de facto* TCP/IP standards, ESnet continues to satisfy the varied international requirements of its user base in addition to providing transition avenues for the applications based on the older proprietary and legacy protocols.

Management

When the new multiprotocol based ESnet came into existence during the latter half of the 1980s it became evident that network management needed to change. ESnet chose to use a pre-commercial version of a Digital Equipment Corporation (DEC) network management tool that handled both the IETF *de facto* standard *Simple Network Management Protocol* (SNMP) and the DECnet IV management protocol. Although the OSI *Common Management Information Protocol* (CMIP) was a standard at the time and being recommended as the government standard by NIST, there were no viable stable products available then, so ESnet used SNMP as the basis for its network management solution and continues to do so today. The ESnet sites also chose to implement SNMP in order to be interoperable with ESnet and to provide an ESnet wide network management framework. That decision was arrived at by consensus, not by mandate. This choice for adoption of a non formal standard based solution has been vindicated, as evidenced by the enormous selection of commercial SNMP products and the number of networks, including government networks, that currently use SNMP as their primary management tool.

Standards—Too much of a good thing? (*continued*)

The *Government Network Management Protocol* (GNMP) is the successor of and a derivative of CMIP and is being championed by NIST as the network management protocol to be used by the federal government. Regardless of technical merit, its fate is yet to be determined. Nevertheless, had ESnet not adopted the *de facto* standard of SNMP as its network management framework it would not have been able to deliver quality operational networking services to its customers during the “shake out time” in the network management arena. This is an example of the successful use of “just in time standards” to achieve interoperability.

New tools

Many times there is no choice but to adopt innovative pre-standard tools and applications in order to address the user's requirements and to provide some base level of interoperability. Information search and retrieval tools, such as *WAIS*, *WWW*, *Mosaic*, and *Gopher* have no formal (OSI) counterpart. ESnet, in the context of providing the richest production oriented environment that meets the needs of its users, has implemented these tools in lieu of waiting for formal standards based or *Federal Information Processing Standards* (FIPS) solutions or alternatives. It is these tools that have captured the imagination of the Clinton Administration and the public when describing the NII. If all agencies activities, including DoE's ESnet, had implemented only the mandated formal GOSIP/OSI standards and the procurement standards of GSA, these enabling technologies would not be available to our researchers and administrators today and the reality of a FII and NII would be much more distant. Thus it is clear that when treating standards as a means, and not as an end onto itself, not only are the chances of addressing the user's requirements enhanced but also those of achieving interoperability.

Directory Service

Systems and standards incorporate the use of versions to allow for evolution, either adding new features or for correcting “bugs.” Even if only one standard or system is chosen there will still be issues relating to interoperability that need to be addressed due to the existence of different versions. One such example is seen with the use of X.500 for providing user based directory services for ESnet and its users. When ESnet implemented its X.500 services in 1990, there were no commercially available and standard based products available. ESnet used the publicly available QUIPU X.500 directory software to provide such services over both the TCP/IP and OSI/GOSIP stacks. The early adoption of this technology has given ESnet users directory services on an operational basis many years ahead of those still waiting for a pure standards based vendor directory services. There are a few interoperability issues between QUIPU and other X.500 services. These are mainly due to the fact that the 1988 version of X.500 was lacking many essential management capabilities, such as replication, access control, and distributed operations. Therefore the QUIPU version had to implement its own version of these capabilities in order to deploy an operational solution before 1992, when these issues were to be addressed by the formal standards process. As QUIPU and other products implement the 1992 X.500 standard these interoperability problems are expected to be ameliorated. Meanwhile, ESnet is using X.500 with NASA, Control Data Corporation, and others. This is a clear example of the standards process taking too long (i.e., the 1988 version did not address necessary operational issues) and also shows that adhering to standards (e.g., 1988 X.500) is not the panacea for interoperability and functionality that many believe it to be.

E-mail ESnet has also implemented an X.400 gateway that provides X.400 and SMTP interoperability for the ESnet community. This early deployment of services has provided the ESnet community and the DoE important X.400 and SMTP e-mail capabilities by pragmatically providing interoperable multi-protocol services instead of needlessly awaiting resolution of a philosophical standards debate. Again, the focus on interoperability and functionality, rather than exclusively choosing one standard, has provided interoperability among the DoE community.

FIX Interoperability is not always achieved through the implementation of a single standard. ESnet peers with other Federal Research and Education (R&E) Entities at the *Federal Internet eXchange* (FIX), where they use inter-network peering protocols such as the *Exterior Gateway Protocol* (EGP) and the *Border Gateway Protocol* (BGP) for exchanging routing and connectivity information and provide for the interconnection and interoperability of federal agency R&E networks. The FIX can be considered a standard for interoperability, at least for the sake of connecting the federal R&E networks [2]. It is important to note that the FIX concept is not bound to any one protocol (standard), it supports DECnet IV, TCP/IP, and OSI/CLNP and others as needed. Even the media used for interconnection, originally Ethernet, is not exclusive. The FIXes have implemented FDDI and can easily provide a hybrid media interconnection point. The success of the FIX concept is seen in other subsequent instances, albeit with added and/or different functionality, such as the *Commercial Internet eXchange* (CIX), *Network Access Points* (NAPs), and the *Global Internet eXchange* (GIX) as exemplified in the Washington D.C. MAE-EAST implementation. We suggest that the success of the FIX and its successors is due to the fact that the focus was on providing a collection of usable standard and non-standard based solutions for achieving interoperability rather than focusing on the identification and exclusive use of a single mandated and/or formal standard.

Video conferencing The effects of an exclusive standard, arrived at through the federal acquisition process policy, can impede interoperability between the government and non-government organizations. In the late 1980s and early 1990s many Energy Research programs were pursuing the use of videoconferencing to augment the current programmatic communications mechanisms. Several sites selected video solutions that were not FTS2000 *Compressed Video Teleconferencing Systems* (CVTS) based. The DoE mandate at that time was to use the GSA Provided Service (a procurement and use standard) that used a non-standards based codec. The ESnet sites had to seek exemption from this mandate in order to acquire a product that could conform to the International Video Standards (H.261 and other appropriate standards). In order to satisfy its videoconferencing and collaborative workspace requirements, the ESnet community developed and is implementing a plan that promotes H.200 and H.300 standard based videoconferencing, dedicated room video solutions, and desktop solutions. The desktop solutions also include workstation based packetized video and audio capabilities across the Internet (e.g., MBONE). Neither the point-to-point standards based videoconferencing systems or packet based desktop systems were a viable solution for exclusively addressing the requirements of the ER community. A plan that incorporated both solutions, including the support of MBONE across ESnet, was the only answer to ER's international and multi-organizational collaborations. OSC is currently funding research that will provide for the interoperability of these videoconferencing systems.

Standards—Too much of a good thing? *(continued)*

Operating systems

The technology life cycle constantly demands the re-evaluation of past decisions and the transition to newer standards and technologies as they evolve. The National Energy Research Supercomputer Center has hosted many serial 1 and “early shipment” supercomputers. In order to provide users with a production environment, they had to design and implement their own operating system (Cray Time Share System—CTSS) and libraries. This was labor intensive, but given the lack of commercial operating systems available for supercomputers at the time, it was the proper choice. In fact, some of the NSF supercomputer centers bootstrapped themselves into operation by using the NERSC and other DoE supercomputer center operating systems and system software. A tough decision presented itself when a commercially available UNIX operating system became available for the supercomputers. The users had become accustomed to features available only in the NERSC developed system and the support staff felt that their system was superior to that of the UNIX based commercial software. However, the rest of the world, including the NSF supercomputer centers, were now using UNIX and the design of POSIX had commenced, thereby creating problems of data and code interoperability between the NERSC and other supercomputer centers. NERSC finally adopted UNIX as its systems level software base in 1992 and went through the painful process of the transition to UNIX. Although the CTSS system may have been superior to the UNIX system in many aspects, the transition should have occurred sooner in order to provide better interoperability between NERSC users and those using other supercomputing facilities, including other agency funded centers. This same scenario is likely to arise again at all sites, not just NERSC, when the transition to Open Software Foundation (OSF) solutions and POSIX compliant systems are enacted, and it will occur again when the move from POSIX or OSF to its successors are made. Evolution demands constant re-evaluation and adaptation.

Emerging technologies

OSC and other areas of DoE have long recognized the value of standards, when they are realistically designed and implemented in a timely fashion, and have directly funded development of necessary standards in addition to participating in standards forums such as the IETF, CCITT (now ITU), the IEEE, the ATM Forum and others. The OSC has funded research in many advanced technology areas, including collaborative work environments and high speed network access. The research and development of high speed networking technologies includes the development of high speed interfaces, such as HIPPI, for use in connecting high speed cycle and storage servers at gigabit speeds, in addition to the investigation of high speed LAN access, such as Fiber Channel, HIPPI, ATM and others, to high speed wide area ATM network services. OSC's support of standards development in niche areas, such as HIPPI or packetized audio, has helped to provide necessary standards and solutions for areas not normally pursued by industry due to their perceived small customer base. It is interesting to note that DoE funds the research and development (and associated standards activities) of diverse and competing solutions, such as HIPPI, ATM, and Fiber Channel, without prejudging or pre-selecting the winners. The end result is that the standards and solutions that have garnered outside support, address real user requirements, and have lived through the natural selection process are chosen and implemented. It also allows for the use of all three technologies, if necessary, without prematurely pruning a branch off of the technology evolution tree.

Security

The federal government should not think of itself as an island when choosing standards and solutions. The use of federally mandated standards in a multi-organizational (i.e., includes more than federal government organizations), widely dispersed, and collaborative environment, such as the Internet, introduces many issues not encountered in an intragovernmental situation. The DoE has been instructed that it must use hardware based DES to encrypt unclassified but sensitive information. The strict adherence to this directive has precluded the use of other security techniques, such as *Privacy Enhanced Mail* (PEM), for the purpose of notifying DoE sites, principle investigators, and collaborators at both DoE and non-DoE facilities of suspected or known viruses, worms, and other attacks. The adherence to this directive impedes DoE site efforts to communicate known security breeches to its non-DoE and non-Governmental collaborators in a timely fashion, thereby increasing the other sites' risk of being compromised. Governmental directives and guidelines, with respect to computer and telecommunications standards, should take into account that government agencies and the FII must be an integral part of the NII, when agencies collaborate with or serve a multitude of non-government entities and organizations. The choice of one exclusive security solution (hardware based DES) has hindered the government's ability to interoperate with the public sector.

Naming and addressing

Many of the above real life examples show that waiting for formal standards and/or implementing one set of mandated standards does not necessarily provide interoperability or an operational system. They have also shown that interoperability is sometimes achieved without exclusively using standards. Success is achieved by implementing a full array of standards and solutions at different maturity levels and at different times on the technology and standards continuum. However, there are many areas that could greatly benefit from standards that are often overlooked. Naming and addressing are a very important part of any infrastructure, yet they are often overlooked or addressed as an afterthought. One of successes of the Internet was the availability of a known name and address space that covered anyone, including the Federal Government and other nations. The OSI naming infrastructure in the US was fragmented, with a set of rules and regulations established and mandated for the federal government but none for the private sector. We believe that the lack of a US wide OSI name and address registrar, complete with standards arrived at by a wide consensus, was another major impediment to the success of OSI in the US. Even today, there is still no US registrar for X.400 management domain names and X.500 top level directory domain servers. A standard for handling and administering OSI names and addresses, including those for operational purposes, is more important for interoperability than the standards used by the mail transfer agents as described in CCITT standards documents. Many users and suppliers of the NII would benefit from an architecture that supports a standard, cheap, and easily used registrar and process for name and address management, further simplifying the user interface to the system.

Procurement standards

Procurement policies and practices, such as FTS2000, also set standards indirectly through the policy of requiring agencies to exclusively acquire and use certain services, standards, hardware, software and systems. When government wide contracts are let, such as FTS2000, they usually result in one or two winners and subsequently reduce the number of alternative, and possibly innovative, technological solutions and competitive service providers available to the government.

Standards—Too much of a good thing? *(continued)*

The “gene pool” of future technologies is enhanced by the availability of multiple solutions and offerings. The raw purchasing power of the federal government, as enacted through a single agency such as GSA, strongly affects, and the authors believe it to be an adverse effect, the competitive marketplace and indirectly establishes standards (e.g., FTS2000 CVTS) by wielding the awesome purchasing club of the federal government. Awardees of large government wide contracts are guaranteed both customers and a profit over a long period of time and therefore have little incentive to introduce new technologies at a rate commensurate with technology evolution (e.g., standards based video-conferencing codecs) nor to provide interoperability (e.g., interoperability between FTS2000 networks A and B and from FTS2000 to the Internet), thereby affecting the agencies’ abilities to interoperate with non-government organizations (e.g., the government was using FTS-2000 X.400 or proprietary services while the rest of the nation used TCP/IP). Future procurement policies, such as Post-FTS2000 [16], need to be defined such that services and standards are chosen from among those that will provide interoperability with the rest of the nation and allow for a multitude of service providers in order to ensure competition and the injection of new and innovative technologies into the government workplace.

Conclusions

As seen in the above examples, the DoE (especially those programs funded and sponsored by the OSC), has been a proponent of standards for the purpose of enhancing interoperability between its users, other agencies, academia, industry, and the public sector on both a national and international basis. Yet its adoption and use of these standards has been driven from a pragmatic perspective that acknowledges the fact that new requirements and solution spaces will be iteratively introduced into the technology and standards continuum and that a rational approach to adoption and use must be employed. This includes: 1) the participation in SDOs and development of new standards when necessary (e.g., HIPPI), 2) the adoption of pre-standards implementations (e.g., MFENET, Packet Video) in lieu of subscribing to the “Waiting for Godot” [2] standards scenario where someone or some organization will wait forever for an event that either takes an inordinate amount of time or will not occur at all, 3) assuming and planning for the co-existence of multiple standards (e.g., TCP/IP and OSI), whether they are different protocols or versions thereof, 4) having the conviction and insight to substitute a new set of standards for an existing one when appropriate (e.g., TCP/IP for MFENET), and 5) adoption of standards based commercial products when they are easily obtainable and satisfy requirements (e.g., multi-protocol routers). In addition, the combined use of multiple standards from multiple SDOs and in different states of maturity has proven successful in addressing the needs and requirements of the user community.

Furthermore, the old standards setting process is obsolete. We agree with Cliff Lynch’s statement that “The era of Standards as an end product has ended.” [9] In addition, we should embrace this technology and standards based (r)evolution and use it as leverage to redefine the method with which the federal government specifies, uses, and relies on standards. The authors also believe that the natural selection process, which is sometimes perceived as random and chaotic, found in the technology and standards continuum is not necessarily a bad condition and that these variables combined with diverse approaches, ideas, and technologies, can ensure a rich intellectual and technological future, and thereby provide the competition necessary for a healthy technology and information based economy.

As is found in the biological world, the natural selection process applied to a wide variety of alternatives is part of the normal evolution cycle and enhances the probability that the best solution(s) survive. If the Internet had not been developed and used we could be currently constrained to using X.25 services. It is hard to imagine how today's information innovations such as multimedia, *World Wide Web*, *Gopher*, and others would have been introduced if the agencies had exclusively implemented OSI. The *Government OSI Profile* (GOSIP) is a procurement profile that references OSI and other protocols to be acquired by Government Agencies. GOSIP does not preclude the use or acquisition of other protocols not specified in the US GOSIP; however, the result has been an interpretation by many Information Resources Management (IRM) officials, at all levels, that since the GOSIP is a procurement specification it was also intended to be the primary choice for implementation, with little room for exceptions. This attitude has resulted in the same effect as if the GOSIP was exclusively mandated for use, with the noted exception of the Research and Education communities within the agencies who have been successfully using the Internet and TCP/IP. This overzealous mandate has impeded the governments ability to interoperate with the non-federal community and will also impede the deployment of the NII. Let's not add our current "dominant" technological gene pool to the endangered species list by subjugating it to exclusive long drawn out formal standards processes.

Pending legislation, S.4—The National Competitiveness Bill, states "federal government contribution of resources and more active participation in the voluntary standards process in the US can increase their compatibility with the standards of other countries, and ease access of products manufactured in US manufacturers to foreign markets" and "the federal government, working in cooperation with private sector organizations including trade associations, engineering societies, technical organizations, and other standards-setting bodies can effectively promote Federal Government use of United States consensus standards and, where appropriate, the adoption and Federal Government use of international standards." The need to compete on an international basis, and therefore use international standards, should not prevent us from continuing our lead role in the technological and standards arenas by forging ahead in those areas where we can continue to be leaders.

Finally, the selection of technologies and standards may have a strong impact on many lives and businesses. Therefore, it is incumbent upon the analysts and federal employees who affect policy [18] in these areas and establish standards setting processes to make sure that they consider the issues in a holistic manner and not focus solely on short term financial and efficiency issues; specifically, they need to understand that they cannot truly benefit the federal government process without taking into consideration other variables and their affect on the future of our nation. As an example, anticipatory standards may prematurely prune a branch from the technology evolution tree and subsequently kill a technology that might prove to be very economical and save the government a lot of money in the future, or be an enabler technology that was necessary for solving national challenges, such as finding a cure for cancer. The process must also resist the temptation to treat standards in a simple fashion where standards are viewed as being only technical or political.

Standards—Too much of a good thing? (*continued*)

Since standards eventually move from paper to implementation, it is imperative that the process involve persons with recent real life operational experience to give them the necessary insight for choosing standards wisely—not just making a politically correct choice. The Government's current practice of setting standards (e.g., Clipper Chip, ADA, GOSIP, FTS2000, etc.) with the intent of creating supportive markets for purposes of saving the government money, providing interoperability, and satisfying security agencies' programs needs to be addressed in context of this moral dilemma since those making the decisions create winners and losers in addition to possibly prematurely affecting the technological gene pool of the future. These are economical, societal and ethical issues in addition to being both political and technical issues.

References

- [1] "20/20 Vision: The Development of a National Information Infrastructure," U.S. Department of Commerce, 1994.
- [2] Aiken, Robert J., Braun, Hans-Werner, Ford, Peter, "NSF Implementation Plan for Interagency Interim NREN," *Journal for High Speed Networks*, Volume 2, No. 1, 1993.
- [3] Becket, Samuel, *Waiting for Godot*, Grove Press, 1954.
- [4] Drucker, Peter F., *Post-Capitalist Society*, Harper-Collins, 1993.
- [5] "Global Standards: Building Blocks for the Future," U.S. Congress, Office of Technology Assessment, 1992.
- [6] "High Performance Computing and Communications: Toward a National Information Infrastructure," Committee on Physical, Mathematical, and Engineering, and Technology, Office of Science and Technology Policy, 1993.
- [7] Information Infrastructure Task Force, "The National Information Infrastructure: Agenda for Action," National Telecommunications and Information Administration, 1993.
- [8] Kuhn, Thomas S., *The Structure of Scientific Revolutions*, The University of Chicago Press, 1962.
- [9] Lynch, Clifford, "Interoperability: The Standards Challenge for the 1990s," *The Wilson Library Bulletin*, 1993.
- [10] "National Performance Review: Reengineering Through Information Technology," Office of the President, 1993.
- [11] National Research Council, "Crossroads of Information Technology Standards," National Academy Press, 1990.
- [12] National Research Council, "Keeping the U.S. Computer Industry Competitive," National Academy Press, 1992.
- [13] Office of Management and Budget, "Government Electronic Mail for the Federal Government," 1994.
- [14] Office of the President of the United States, "Streamlining Procurement Through Electronic Commerce," Executive Memorandum, 1993.
- [15] "Perspectives on the National Information Infrastructure: Ensuring Interoperability," Computer Systems Policy Project, 1994.

- [16] "Post-FTS2000 Acquisition Alternatives White Paper," Acquisition Working Group of the Interagency Management Council, 1994.
- [17] Tasse, Greg, "The Roles of Standards as Technology Infrastructure," National Institute of Standards and Technology, 1993
- [18] Tong, Rosemarie, *Ethics in Policy Analysis*, Prentice Hall, 1986.

ROBERT J. AIKEN possesses a Masters Degree in Computer Science from Temple University. He is currently on "detail" from Lawrence Livermore National Laboratories to the Department of Energy's Energy Research Office of Scientific Computing where he consults on network research, telecommuting, energy demand management systems, NREN, ESnet, High Performance Computing and Communications (HPCC), and the National Information Infrastructure (NII). Bob also served as the NREN program manager at the National Science Foundation where he coauthored the next generation NSFNET and the Network Access Points (NAPS). Prior to Bob's DC foray he created and managed the ESnet Network Information Services/E-mail/Directory group. Bob was an Assistant Professor of Computer Science at Hood College and has worked on supercomputers, very high speed LANs, and distributed systems at Sandia National Laboratories. Bob is easily recognized by his trademark cowboy boots and eclectic tie collection. He is a member of the ACM, IEEE, and ISOC and can be reached at aiken@es.net

JOHN CAVALLINI holds a BA in Mathematics from Wilkes College and an MS in Electrical Engineering majoring in Computer Science from George Washington University. John has worked in operating systems and performance optimization for large scale systems at NASA Goddard Space Flight Center. He has managed large Federal Information Resources Management programs at the Health and Human Services Department and at the National Oceanic and Atmospheric Administration. John now heads up the Department of Energy's Scientific Computing Office where he initiated the Department's supercomputer access program, the Energy Sciences Network program, and most recently its High Performance Computing and Communications program. John can be reached at cavallini@nersc.gov

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use this address for letters to the Editor, questions about back issues etc.:

ConneXions—The Interoperability Report

303 Vintage Park Drive

Suite 201

Foster City, CA 94404-1138


USA

Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)

Fax: +1 415-525-0194

E-mail: connexions@interop.com

For questions about your subscription please call our customer service hotline: 1-800-575-5717, (+1 502-493-3217 outside the US). This is the number for our subscription agency, the Cobb Group. Their fax number is +1 502-491-8050. The mailing address for subscription payments is P.O. Box 35840, Louisville, KY 40232-9496.



September 12-16, 1994

To register call:

1-800-488-2883

or

1-415-578-6900

This publication is distributed on an "as is" basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *ConneXions—The Interoperability Report*.

Call for Papers

The *Internet Society Symposium on Network and Distributed System Security* will be held February 16–17, 1995 at the Catamaran Hotel in San Diego, California.

Goal The symposium will bring together people who are building software and/or hardware to provide network and distributed system security services. The symposium is intended for those interested in the more practical aspects of network and distributed system security, focusing on actual system design and implementation, rather than in theory. We hope to foster the exchange of technical information that will encourage and enable the Internet community to apply, deploy and advance the state of the available security technology. Symposium proceedings will be published by the Internet Society.

Topics Topics for the symposium include, but are not limited to, the following:

- Design and implementation of security services—access control, authentication, availability, confidentiality, integrity, and non-repudiation.
- Design and implementation of security mechanisms and support services—encipherment, authentication, and key management systems, including fair cryptography—access control, authorization and audit systems—and intrusion detection systems.
- Requirements and designs for securing distributed applications and network functions—message handling, file transport, remote file access, directories, time synchronization, interactive sessions, remote data base management and access, routing, voice and video multicast and conferencing, news groups, network management, boot services, mobile computing, and remote I/O.
- Requirements and designs for securing networked information resources and tools—*Archie* servers, the *Wide Area Information Servers* (WAIS), the *Internet Gopher*, and the *World Wide Web* (WWW).
- Design and implementation of measures for controlling inter-network communication and services in a coherent manner—firewalls, packet filters, application gateways, and user/host authentication schemes.
- Requirements and designs for telecommunications security especially for emerging technologies—very large systems like the international Internet, high-speed systems like the gigabit test-beds, wireless systems, and personal communication systems.
- Special issues and problems in security architecture, such as interplay between security goals and other goals—efficiency, reliability, interoperability, resource sharing, and cost.

General Chair

Jim Ellis

CERT Coordination Center

Program Chairs

David Balenson
Rob Shirey

Trusted Information Systems
The MITRE Corporation

Program Committee

Tom Berson,	Anagram Laboratories
Matt Bishop,	University of California at Davis
Ravi Ganesan	Bell Atlantic
Burt Kaliski	RSA Laboratories
Steve Kent	BBN Communications
Paul Lambert	Motorola
John Linn	OpenVision Technologies
Clifford Neuman	Information Sciences Institute
Hilarie Orman	University of Arizona
Michael Roe	Cambridge University (UK)
Rob Rosenthal	NIST
Jeff Schiller	Massachusetts Institute of Technology
Mike St. Johns	Advanced Research Projects Agency
Peter Yee	NASA
Roberto Zamparo	Telia Research (Sweden)

Submissions

The committee invites both technical papers and proposals for panel discussions on technical and other topics of general interest. Technical papers should be 10–20 pages in length. Panel proposals should be two pages in length, and should describe the panel topic, name the panel chair, explain the format of the panel, and list three to four potential panelists. The technical papers will appear in the proceedings. Panel chairs and panelists will have the option of having written statements appear in the proceedings.

All submissions should contain a separate title page which includes the type of submission (paper or panel), the title or topic, the names of the author(s), organizational affiliation(s), telephone and fax numbers, postal addresses, Internet electronic mail addresses, and the point of contact, if more than one author. Since the review process will be anonymous, the author's names, affiliations and other information should appear only on the separate title page.

Submissions should be made via electronic mail. Submissions may be in either of two formats: *PostScript* or ASCII. If the committee is unable to print a *PostScript* submission, it will be returned and ASCII requested. Therefore, *PostScript* submissions should arrive well before 15 August. If electronic submission is absolutely impossible, submissions should be sent via postal mail. All submissions and other correspondence should be directed to the Program Co-Chair:

David M. Balenson
Trusted Information Systems, Inc.
3060 Washington Road (Rt. 97)
Glenwood, Maryland 21738 USA
Phone: +1 301 854-6889
FAX: +1 301 854-5363
E-mail: balenson@tis.com

Each submission will be acknowledged through the medium by which it is received. If acknowledgment is not received within seven days, please contact the Program Co-Chair as indicated above. Instructions for preparing camera-ready copy for the proceedings will be postal mailed at the time the papers are accepted.

Important dates

Deadline for paper submission:	August 15, 1994
Notification sent to authors by:	October 17, 1994
Deadline for camera-ready copy:	November 15, 1994

CONNEXIONS

303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf
Senior Vice President, MCI Telecommunications
President, The Internet Society

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to **CONNEXIONS**).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

303 Vintage Park Drive, Suite 201
Foster City, CA 94404-1138
415-578-6900 FAX: 415-525-0194
connexions@interop.com

Back issues available upon request \$15./each
Volume discounts available upon request

CONNEXIONS